arXiv:1405.0237v1 [cs.IT] 1 May 2014

An RIP-based approach to $\Sigma\Delta$ quantization for compressed sensing

Joe-Mei Feng and Felix Krahmer

May 2, 2014

Abstract

In this paper, we provide a new approach to estimating the error of reconstruction from $\Sigma\Delta$ quantized compressed sensing measurements. Our method is based on the restricted isometry property (RIP) of a certain projection of the measurement matrix. Our result yields simple proofs and a slight generalization of the best-known reconstruction error bounds for Gaussian and subgaussian measurement matrices.

1 Introduction

1.1 Compressed sensing

Compressed sensing has drawn significant attention since the seminal works by Candès, Romberg, Tao [8], and Donoho [14]. The theory of compressed sensing is based on the observation that various cases of natural signals are approximately sparse with respect to certain bases or frames. The basic idea is to recover such signals from a small number of linear measurements. Hence the problem turns into an underdetermined linear system. Various criteria have been proposed to determine whether such a system has a unique sparse solution. In this paper we will work with the restricted isometry property (RIP) as introduced by Candès et al. [9] in the context of recovery guarantees for ℓ_1 minimization.

Definition 1. A matrix $A \in \mathbb{R}^{m\times N}$ has the restricted isometry property (RIP) of order s if there exists $0 < \delta < 1$ such that for all s-sparse vectors $x \in \mathbb{R}^N$, i.e., vectors that have at most s non-zero components, one has

$$(1-\delta)\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1+\delta)\|x\|_2^2.$$

The smallest such δ is called the restricted isometry constant of order s and is denoted by δ_s.

There have been a number of works on recovery guarantees for compressed sensing with RIP measurement matrices. Recovery can be guaranteed for various algorithms. For the original context of ℓ_1 minimization, the most recent

results require the measurement matrix to have a restricted isometry constant of $\delta_{2s} < \frac{1}{\sqrt{2}}$ [7], which is known to be optimal [11].

Finding the restricted isometry constant of a measurement matrix is, in general, an NP hard problem [31]. On the other hand, deterministic matrix constructions with guaranteed RIP are only known for relatively large embedding dimensions (see for example [13]). That is why many papers on the subject work with random matrices.

Examples of random matrices known to have the RIP for large enough embedding dimension with high probability include subgaussian, partial random circulant [23], and partial random Fourier matrices [29]. A subgaussian matrix has independent random entries whose tails are dominated by a Gaussian random variable (cf. Definition 2). Such matrices have been shown to have the RIP provided $m = \Omega(s \log(eN/s))$, see for example [2]. This order of the embedding dimension m is known to be optimal [15]. Examples of subgaussian matrices include Gaussian and Bernoulli matrices.

1.2 Quantization

To allow for digital transmission and storage of compressed sensing measurements, one needs to quantize these measurements. That is, the measurements need to be represented by finitely many symbols from a finite alphabet. In this paper, we only consider alphabets consisting of equispaced real numbers. The extreme case of considering the set of only the two elements $\{-1, 1\}$ is also called 1-bit quantization.

The most intuitive method to quantize the measurements is to map each of them to the closest element from the alphabet. Since this method processes the quantization independently for each measurement, it is also called memoryless scalar quantization (MSQ).

Most of the literature on MSQ compressed sensing up to date considers 1-bit quantization [6, 22, 28, 1], which amounts to considering only the measurement signs. Jacques et al. [22] showed that for Gaussian measurements or measurements drawn uniformly from the unit sphere, a reconstruction error of $O(\frac{s}{m} \log \frac{mN}{s})$ is feasible. However, they did not provide an efficient algorithm that guarantees this accuracy. Later, for Gaussian measurements, Gupta et al. [18] demonstrated that one may tractably recover the support of a signal from $O(s \log N)$ measurements. Plan et al. [28] showed that one can, again for Gaussian measurements, reconstruct the direction of an s-sparse signal via convex optimization, with accuracy $O((\frac{s}{m})^{\frac{1}{5}})$ up to logarithmic factors with high probability. Ai et al. [1] derived similar results for subgaussian measurements under additional assumptions on the size of the signal entries.

On the other hand, in [22] it was shown that the ℓ_2 reconstruction error can never be better than $\Omega(\frac{s}{m})$. To break this bottleneck of MSQ, $\Sigma\Delta$ quantization for compressed sensing has drawn attention recently. $\Sigma\Delta$ quantizes a vector as a whole rather than the components individually, i.e., the quantized values depend on previous quantization steps.

$\Sigma\Delta$ quantization was originally introduced as an efficient quantizer for redundant representation of oversampled band-limited functions [21]. Later on, a rigorous mathematical error analysis was provided by [10] and many follow-up papers. The best known error decay rates are exponential in the oversampling rate, as derived in [16, 12]. This is known to be optimal: in [26], corresponding lower bounds are derived, which also show that the achievable accuracy must depend on the signal amplitude.

In [3], $\Sigma\Delta$ has been extended to frame expansions; this will be also the viewpoint taken in this paper. The first works on $\Sigma\Delta$ schemes for frame quantization, such as [3, 5], required frame constructions with particular smoothness properties to yield reconstruction guarantees. In [27], the authors observed that what is needed is in fact a requirement on the dual frame used for reconstruction rather than the frame itself. Reconstruction guarantees can hence be improved by choosing the dual frame used for reconstruction appropriately. Optimizing the dual frame in this respect led to the definition of Sobolev dual frames [4], cf. Section 2.2.2 below. Combined with the exponential error bounds derived for the corresponding $\Sigma\Delta$ schemes for bandlimited functions [16, 12], Sobolev dual reconstructions yield root-exponential error decay in the oversampling rate. This constitutes the best known accuracy guarantees for coarse frame quantization, both for harmonic frames and Sobolev self-dual frames [24] and for subgaussian random frames [25].

Sobolev dual reconstructions have also been crucial for being able to apply $\Sigma\Delta$ quantization to compressed sensing measurements. Güntürk et al. [17] proved the first recovery guarantees for this setup, showing that for rth order $\Sigma\Delta$ quantization applied to Gaussian compressed sensing measurements, the ℓ_2 reconstruction error is of order $O((\frac{s}{m})^{\alpha(r-\frac{1}{2})})$ with high probability. Here $\alpha \in (0,1)$ is a parameter and the required measurements grows with α, tending to infinity as $\alpha \to 1$. Indeed for r large enough this breaks the MSQ bottleneck. More recently, in [25], this result has been generalized to subgaussian measurements.

1.3 Contributions

The main contribution of this paper is that the restricted isometry property (RIP) is applied to estimate the error bound for $\Sigma\Delta$ quantized compressed sensing. That is, once we know the restricted isometry constant of a modification of the measurement matrix, we can estimate the reconstruction error.

In the following results, we assume that the $\Sigma\Delta$ quantized measurements with quantization alphabet $\mathcal{Z} = \Delta\mathbb{Z}$, $\Delta > 0$, are given. We refer the readers to Section 2.2 for details on the quantization scheme employed. A special role is played by the rth power of the inverse of the finite difference matrix D as introduced in (2) below; denoting the singular value decomposition of D^{-r} by $D^{-r} = U_{D^{-r}} S_{D^{-r}} V^*_{D^{-r}}$, we obtain our main theorem given as follows.

Theorem 1. *Suppose one is given a measurement matrix $\Phi \in \mathbb{R}^{m\times N}$ such that both Φ and $\sqrt{\frac{1}{\ell}} P_\ell V^*_{D^{-r}} \Phi$, $\ell \leq m$ have the restricted isometry constant $\delta_{2s} < \frac{1}{\sqrt{2}}$,*

where P_ℓ maps a vector to its first ℓ components.

Then for an s-sparse signal $x \in \mathbb{R}^N$ satisfying $\min_j |x_j| \geq K2^{r-\frac{1}{2}}\Delta$, for some positive constant K, denote by q the rth order $\Sigma\Delta$ quantized measurements of Φx with step size Δ. Furthermore, denote by T the support set recovered from Φx via ℓ_1 minimization and choose $L_{sob,r}$ to be the Sobolev dual matrix of Φ_T (see Section 2.2.2 for details). Then reconstructing the signal via $\hat{x}_T = L_{sob,r}q$ yields a reconstruction error bounded by

$$\|x - \hat{x}\|_2 \leq C\Delta(\frac{m}{\ell})^{-r+\frac{1}{2}},$$

where $C > 0$ is a constant depending only on r.

Note from Theorem 1 that smaller values of ℓ yield better error bounds. However, ℓ has to be large enough such that $\frac{1}{\sqrt{\ell}}(P_\ell V_{D^{-r}}^* \Phi)$ has the restricted isometry constant $\delta_{2s} \leq \frac{1}{\sqrt{2}}$.

This result can be applied to obtain recovery guarantees for Gaussian and subgaussian measurements (in the sense of Definition 2 below). The resulting bounds for the first two cases agrees with those derived in [17] and [25], as summarized in the following Theorem.

Theorem 2 ([17, 25])**.** *Let Φ be an $m \times N$ matrix whose entries are independent, mean zero, unit variance ρ-subgaussian random variables and suppose that $\lambda := m/k \geq (C\log(eN/k))^{\frac{1}{1-\alpha}}$ where $\alpha \in (0,1)$. With high probability the rth order $\Sigma\Delta$ reconstruction $\hat{x}$ satisfies*

$$\|x - \hat{x}\|_2 \leq C'\lambda^{-\alpha(r-1/2)}\delta,$$

for all $x \in \Sigma_k^N$ for which $\min_{j\in \mathrm{supp}(x)} |x_j| > K'\Delta$. Again, Δ is the step size of the $\Sigma\Delta$ quantization alphabet and C, C', K' are appropriate constants that depend only on r and ρ.

1.4 Organization

The paper is organized as follows. We first introduce in Section 2 some background and previous results on $\Sigma\Delta$ quantization, suprema of chaos processes, and the partial random circulant matrices. In Section 3 we present our main result showing how the RIP is used to estimate the reconstruction error for quantized compressed sensing. In Section 4, we explain how our result recovers the best-known bounds for Gausssian and subgaussian measurement matrices using a simple argument, in Section 5 we slightly generalize these bounds. We conclude in Section 6.

2 Background and previous results

2.1 Notation

Throughout this paper, we use the following notation. The set $D_{s,N} = \{x \in \mathbb{R} | \|x\|_2 \leq 1, \|x\|_0 \leq s\}$ is the set of unit norm s-sparse vectors. The ℓ_0-norm $\|\cdot\|_0$

counts the number of non-zero components of a vector. Given a signal x, the support set of x, in short, supp x, is the index set of the non-zero components. The ℓ_2-operator norm is denoted by $\|A\|_{2\to 2} = \sup_{\|x\|_2=1} \|Ax\|_2$. For a matrix A, $\sigma_i(A)$ and $\sigma_{\min}(A)$ denote the ith largest and the smallest singular value, respectively. Furthermore we write $\gtrsim$ and $\lesssim$ to denote $\geq$ or $\leq$ up to a positive multiplicative constant. The Moore-Penrose pseudoinverse of a matrix A is denoted by $A^\dagger = (A^*A)^{-1}A^*$.

We will mainly study subgaussian random matrices, that is, matrices with independent subgaussian entries in the sense of the following definition.

Definition 2. A random variable X is called ρ-subgaussian if $\mathbb{P}(|X| \geq t) \leq 2\exp(-t^2/2\rho^2)$.

2.2 $\Sigma\Delta$ Quantization

In this paper, we exclusively focus on quantization alphabets $\mathcal{Z}$ such that $\mathcal{Z} = \Delta\mathbb{Z}$, for some $\Delta > 0$. Note that while this is an infinite set, one can show that in fact only a finite range of values are assumed [17, 24]. Hence this setup is in line with requiring a finite alphabet. The idea of rth order $\Sigma\Delta$ quantization is to quantize each component of a vector taking the previous r quantization steps into account. More explicitly, a greedy rth order $\Sigma\Delta$ quantization scheme maps a sequence of inputs (y_j) to elements $q_i \in \mathcal{Z}$ via an internal state variable u_i chosen to satisfy the recurrence relation

$$(\Delta^r u)_i := \sum_{j=0}^{r} \binom{r}{j} (-1)^j u_{i-j} = y_i - q_i, \tag{1}$$

where q_i is chosen such that $|u_i|$ is minimized (Note that only in this equation, Δ denotes the finite difference operator, whereas all other occurences in this paper refer to the quantization step size).

With the initial condition $(u_i)_{i=0}^{-\infty} = 0$, Equation (1) can be expressed as

$$D^r u = y - q,$$

where the finite difference matrix $D \in \mathbb{R}^{m\times m}$ is given by

$$D_{ij} \equiv \begin{cases} 1 & , \text{ if } i = j, \\ -1 & , \text{ if } i = j+1, \\ 0 & , \text{ otherwise.} \end{cases} \tag{2}$$

2.2.1 Support set recovery

Given an s-sparse signal x, and an $m\times N$ measurement matrix Φ, where $m \ll N$, we acquire measurements $y = \Phi x$. Applying an rth order $\Sigma\Delta$ quantization scheme to y, we obtain q. Treating q as perturbed measurements, i.e., $q = y + e = \Phi x + e$, one can determine the support set. This is a consequence of the following observation, which is a modified version of Proposition 4.1 in [17] combined with the reconstruction guarantees in [7].

Proposition 1. *Given $\epsilon > 0$ as well as $x \in \mathbb{R}^N$ an s-sparse signal with $\operatorname{supp} x = T$ and $\min_{j\in T} |x_j| \geq K\frac{\epsilon}{\sqrt{m}}$. Here K is an absolute constant. Let $\Phi \in \mathbb{R}^{N\times m}$ be a measurement matrix such that $\frac{1}{\sqrt{m}}\Phi$ has the RIP with $\delta_{2s} < \frac{1}{\sqrt{2}}$. Denote by $e \in \mathbb{R}^m$ a noise vector with $\|e\|_2 \leq \epsilon$, and let x' be the signal reconstructed from the noisy measurements $q = \Phi x + e$ via ℓ_1 minimization, i.e.,*

$$x' = \arg\min \|z\|_1 \text{ subject to } \|\Phi z - q\|_2 \leq \epsilon.$$

Then the index set of largest s components of x' is T, that is, the support set of x is correctly recovered.

Note that in this result, the measurement matrix Φ is not normalized, while in the compressed sensing literature, it is common to normalize the measurement matrix such that it has unit-norm columns. This is because for normalized matrix columns, each measurement will be of order $\frac{1}{\sqrt{m}}$, so quantizing it with a fixed step size Δ will lead to worse and worse resolution. To allow for a fair comparison when m grows, the measurements should rather be chosen independently of m. Therefore, in this paper as well as in [17] the measurement matrices are not normalized, each entry of the measurement matrices is chosen to have variance one.

To apply Proposition 1 to greedy $\Sigma\Delta$ quantization, one sets $e = q - y$, where q is the quantized measurement vector. Elementary estimates (cf. [17]) yield that $\|q - y\|_2 \leq 2^{r-1}\Delta\sqrt{m}$. Thus one obtains that ℓ_1 minimization recovers the correct support set provided that $\frac{1}{\sqrt{m}}\Phi$ has restricted isometry constant $\delta_{2s} < \frac{1}{\sqrt{2}}$ and $\min_j |x_j| \geq K2^{r-\frac{1}{2}}\Delta$.

2.2.2 Estimating the error and the Sobolev dual

When the support set T has been identified, we solve for x using some left inverse of Φ_T, say L. Then the reconstruction ℓ_2-error is given by

$$\begin{aligned}\|x - \hat{x}\|_2 &= \|Ly - Lq\|_2 = \|L(y - q)\|_2 \\ &= \|L(D^r u)\|_2 \leq \|LD^r\|_{2\to 2}\|u\|_2.\end{aligned}$$

The Sobolev dual matrix $L_{sob,r}$, first introduced in [4], is a left inverse of Φ_T defined to minimize $\|LD^r\|_{2\to 2}$, i.e.,

$$L_{sob,r} = \arg\min_L \|LD^r\|_{2\to 2} \quad \text{subject to } L\Phi_T = I.$$

The geometric intuition is that this dual frame is smoothly varying.

As in [17], the explicit formula $L_{sob,r}D^r = (D^{-r}\Phi_T)^\dagger$ yields the error bound

$$\begin{aligned}\|x - \hat{x}\|_2 &\leq \|(D^{-r}\Phi_T)^\dagger\|_{2\to 2}\|u\|_2 \\ &= \frac{1}{\sigma_{min}(D^{-r}\Phi_T)}\|u\|_2 \leq \frac{\Delta\sqrt{m}}{2\sigma_{min}(D^{-r}\Phi_T)},\end{aligned} \tag{3}$$

where the last inequality is derived in [17].

A key ingredient to bounding $\sigma_{\min}(D^{-r}\Phi_T)$ is the following result from the study of Toeplitz matrices, which depends heavily on Weyl's inequality [20] (see for example [17]).

Proposition 2. *Let r be any positive integer and D be as in (2). There are positive constants $c_{s_1}(r)$ and $c_{s_2}(r)$, independent of m, such that*

$$c_{s_1}(r)(\frac{m}{j})^r \leq \sigma_j(D^{-r}) \leq c_{s_2}(r)(\frac{m}{j})^r,\ j = 1, \ldots, m.$$

3 RIP-based error analysis

In this section we will give the quantized compressed sensing problem a mathematical model, and explain how we approach the reconstruction error via the RIP. In the next two sections we show its applications. From Section 2.2.2, the main issue to estimate the reconstruction error is to estimate $\sigma_{min}(D^{-r}\Phi_T)$. Finding the supremum of this expression over all potential support sets T can be interpreted as finding the supremum of the smallest image under $D^{-r}\Phi$ over all unit norm s-sparse vectors. This motivates the connection to the RIP.

In the following proof we show how the RIP can be applied to find this effective smallest singular value.

Proof of Theorem 1. As the assumptions of the theorem are stronger than those of Proposition 1, we conclude that the support is correctly recovered. Based on this observagtion, we now show the error bound. Recall that $D^{-r} = U_{D^{-r}} S_{D^{-r}} V^*_{D^{-r}}$. Then, as S is a diagonal matrix,

$$\begin{aligned}
\sigma_{\min}(D^{-r}\Phi_T) &= \sigma_{\min}(S_{D^{-r}} V^*_{D^{-r}}\Phi_T) \\
&\geq \sigma_{\min}(P_\ell S_{D^{-r}} V^*_{D^{-r}}\Phi_T) \\
&= \sigma_{\min}((P_\ell S_{D^{-r}} P^*_\ell)(P_\ell V^*_{D^{-r}}\Phi_T)) \\
&\geq s_\ell \sigma_{\min}(P_\ell V^*_{D^{-r}}\Phi_T) \\
&\gtrsim (\frac{m}{\ell})^r \sigma_{\min}(P_\ell V^*_{D^{-r}}\Phi_T), \qquad (4)
\end{aligned}$$

where the final inequality follows from Proposition 2.

Thus we need to bound $\sigma_{\min}(P_\ell V^*_{D^{-r}}\Phi_T)$ uniformly over all possible support sets T. Indeed by the RIP assumption for $\frac{1}{\sqrt{\ell}} P_\ell V^*_{D^{-r}}\Phi$, we obtain that $\sigma_{\min}(P_\ell V^*_{D^{-r}}\Phi_T)$ is uniformly bounded from below by

$$\sqrt{\ell}\sqrt{1 - \tfrac{1}{\sqrt{2}}}. \qquad (5)$$

The theorem follows by combining (3), (4), and (5).

□

4 Gaussian and subgaussian matrices

To illustrate the simplicity of our method, we first present a proof of Theorem 2 for standard Gaussian matrices, i.e. matrices with independent entries $\Phi_{i,j} \sim \mathcal{N}(0,1)$.

Proof of Theorem 2 for Gaussian matrices. Set $\ell := m(\frac{s}{m})^\alpha$. As the second factor is less than 1, one always has $1 \leq \ell \leq m$. Since Φ is a standard Gaussian random matrix, due to rotation invariance $\widetilde{\Phi} := (P_\ell V_{D^{-r}}^* \Phi)$ is also a standard Gaussian random matrix. The assumption on λ implies that its embedding dimension satisfies $\ell \geq Cs\log(\frac{eN}{k})$, so standard results (see, e.g., [2]) yield that for C and C' large enough, both $\frac{1}{\sqrt{\ell}}\widetilde{\Phi}$ and $\frac{1}{\sqrt{m}}\Phi$ have the RIP with constant $\delta_{2s} \leq \frac{1}{\sqrt{2}}$. Applying Theorem 1 (choose $K' = K2^{r-\frac{1}{2}}$), we obtain

$$\|x - \hat{x}\|_2 \lesssim \Delta(\frac{m}{s})^{-\alpha(r-\frac{1}{2})},$$

again with high probability, as desired. □

Sketch of proof of Theorem 2 for subgaussian matrices:
The proof proceeds long the same lines as for Gaussian matrices, except that one cannot use the rotation invariance. To bound the RIP constant of $\frac{1}{\sqrt{\ell}}P_\ell V_{D^{-r}}^* \Phi$, we note that for any $x \in D_{s,N}$, $\|P_\ell V_{D^{-r}}^* \Phi x\|_2^2$ is a quadratic form in the "vectorization" of Φ. Hence its tail decay can be estimated via the Hanson-Wright inequality [19, 30]. The RIP then follows via a union bound over an ϵ-net of $D_{s,N}$. This approach is related to certain steps in the original proof in [25].

Remark 1. Note that a complete proof for $\Sigma\Delta$ recovery guarantees needs both support set recovery and fine recovery (cf. [17, 25]) and in this paper we omitted the details of the former. We argue, however, that this coarse recovery step is straightforwardly based on standard compressed sensing results. Hence the core of our error estimate is really just captured in a few lines.

5 Generalization

In contrast to the techniques presented in [25], our method generalizes to certain random matrices with independent subgaussian columns, but no entrywise independence. As an additional criterion, one needs a type of small ball condition for $P_\ell V_{D^{-r}}^*$ applied to one of the random columns of Φ, (which denoted by Φ_j in the following). That is, one needs to exclude that $\|P_\ell V_{D^{-r}} \Phi_j\|_2$ is small with too large probability. If such a condition holds, the necessary RIP bound follows from a modified version of the RIP bound for matrices with independent subgaussian columns [32]. While we do not consider this to be an important generalization (which is why we refrain from presenting the details), we still believe it shows that our method is stronger than previous approaches, so we see the potential to apply it to more relevant, structured measurement scenarios such as partial random Fourier matrices, partial random circulant matrices, etc.

6 Conclusion

In this work we provided a new technique for bounding the reconstruction error arising in $\Sigma\Delta$ quantization for compressed sensing. In addition to greatly simplifying the proofs for the best known recovery guarantees, the new viewpoint hopefully opens the possibility to study broader classes of measurement matries.

Acknowledgment

The authors would like to thank Rayan Saab for providing helpful comments. This work was supported by the DFG Research Training Group 1023.

References

[1] A. Ai, A. Lapanowski, Y. Plan, and R. Vershynin. One-bit compressed sensing with non-Gaussian measurements. *Lin. Alg. Appl.*, 441:222–239, 2014.

[2] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin. A simple proof of the restricted isometry property for random matrices. *Constr. Approx.*, 28(3):253–263, 2008.

[3] J. J. Benedetto, A. M. Powell, and Ö. Yılmaz. Sigma-Delta ($\Sigma\Delta$) quantization and finite frames. *IEEE Trans. Inform. Theory*, 52(5):1990–2005, 2006.

[4] J. Blum, M. Lammers, A. M. Powell, and Ö. Yılmaz. Sobolev duals in frame theory and Sigma-Delta quantization. *J. Fourier Anal. Appl.*, 16(3):365–381, 2010.

[5] B.G. Bodmann, V.I. Paulsen, and S.A. Abdulbaki. Smooth frame-path termination for higher order sigma-delta quantization. *J. Fourier Anal. Appl.*, 13(3):285–307, 2007.

[6] P. T. Boufounos and R. G. Baraniuk. 1-bit compressive sensing. In *Proc. 42nd Annu. Conf. Inf. Sci. Syst. (CISS) , Princeton, NJ*, pages 16–21. IEEE, Mar. 2008.

[7] T. Cai and A. Zhang. Sparse representation of a polytope and recovery of sparse signals and low-rank matrices. *IEEE Trans. Inform. Theory*, to appear.

[8] E. J. Candès, J. K. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59(8):1207–1223, 2006.

[9] E. J. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. Inform. Theory*, 51(12):4203–4215, 2005.

[10] I. Daubechies and R. DeVore. Approximating a bandlimited function using very coarsely quantized data: A family of stable sigma-delta modulators of arbitrary order. *Ann. Math.*, 158(2):679–710, 2003.

[11] M. E. Davies and R. Gribonval. Restricted isometry constants where ℓ^p sparse recovery can fail for $0 < p \leq 1$. *IEEE Trans. Inform. Theory*, 55(5):2203–2214, 2009.

[12] P. Deift, F. Krahmer, and C. S. Güntürk. An optimal family of exponentially accurate one-bit sigma-delta quantization schemes. *Comm. Pure Appl. Math.*, 64(7):883–919, 2011.

[13] R. DeVore. Deterministic constructions of compressed sensing matrices. *J. Complexity*, 23:918–925, 2007.

[14] D. L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52(4):1289–1306, 2006.

[15] S. Foucart, A. Pajor, H. Rauhut, and T. Ullrich. The gelfand widths of ℓ_p-balls for $0 < p \geq 1$. *J. Complexity*, 26(6):629–640, 2010.

[16] C. S. Güntürk. One-bit sigma-delta quantization with exponential accuracy. *Comm. Pure Appl. Math.*, 56(11):1608–1630, 2003.

[17] C. S. Güntürk, M. Lammers, A. M. Powell, R. Saab, and Ö. Yılmaz. Sobolev duals for random frames and $\Sigma\Delta$ quantization of compressed sensing measurements. *Found. Comput, Math.*, 13(1):1–36, 2013.

[18] A. Gupta, R. Nowak, and B. Recht. Sample complexity for 1-bit compressed sensing and sparse classification. In *Proc. IEEE Int. Symp. Inf. Theory (ISIT) , Austin, TX*, pages 1553–1557. IEEE, Jun. 2010.

[19] D. L. Hanson and F. T. Wright. A bound on tail probabilities for quadratic forms in independent random variables. *Annals Math. Statistics*, 42(3):1079–1083, 1971.

[20] R. A. Horn and C. R. Johnson. *Matrix analysis*. Cambridge university press, 2012.

[21] H. Inose, Y. Yasuda, and J. Murakami. A telemetering system by code modulation-Δ-Σ modulation. *IRE Trans. Space El. Tel.*, (3):204–209, 1962.

[22] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk. Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors. *IEEE Trans. Inform. Theory*, 59(4):2082–2102, 2013.

[23] F. Krahmer, S. Mendelson, and H. Rauhut. Suprema of chaos processes and the restricted isometry property. *Comm. Pure Appl. Math.*, to appear.

[24] F. Krahmer, R. Saab, and R. Ward. Root-exponential accuracy for coarse quantization of finite frame expansions. *IEEE Trans. Inform. Theory*, 58(2):1069–1079, 2012.

[25] F. Krahmer, R. Saab, and Ö. Yılmaz. Sigma-Delta quantization of sub-gaussian frame expansions and its application to compressed sensing. *Inform. Inference*, 3(1):40–58, 2014.

[26] F. Krahmer and R. Ward. Lower bounds for the error decay incurred by coarse quantization schemes. *Appl. Comput. Harmonic Anal.*, 32(1):131–138, 2012.

[27] M. Lammers, A.M. Powell, and Ö. Yılmaz. Alternative dual frames for digital-to-analog conversion in sigma–delta quantization. *Adv. Comput. Math.*, pages 1–30, 2008.

[28] Y. Plan and R. Vershynin. One-bit compressed sensing by linear programming. *Comm. Pure Appl. Math.*, 66(8):1163–1333, 2013.

[29] M. Rudelson and R. Vershynin. On sparse reconstruction from fourier and gaussian measurements. *Comm. Pure Appl. Math.*, 61(8):1025–1045, 2008.

[30] M. Rudelson and R. Vershynin. Hanson-wright inequality and sub-gaussian concentration. *Electron. Comm. Probab.*, 18:1–9, 2013.

[31] A. M. Tillmann and M. E. Pfetsch. The computational complexity of the restricted isometry property, the nullspace property, and related concepts in compressed sensing. *IEEE Trans. Inform. Theory*, 2012.

[32] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. In Y.C. Eldar and G. Kutyniok, editors, *Compressed Sensing: Theory and Applications*, pages 210–268. Cambridge Univ. Press, Cambridge, 2012.